MNRAS **000**, 000–000 (0000) Preprint 6 March 2017 Compiled using MNRAS LATEX style file v3.0

Robust detection of quasi-periodic variability: A HAWKI mini survey of late T dwarfs

S. P. Littlefair[1], B. Burningham[2], Ch. Helling[3]
[1] *Dept of Physics and Astronomy, University of Sheffield, Sheffield, S3 7RH, UK*
[2] *Centre for Astrophysics Research, Science and Technology Research Institute, University of Hertfordshire, Hatfield AL10 9AB, UK*
[3] *SUPA, School of Physics and Astronomy, University of St Andrews, St Andrews KY16 9AA, UK*

arXiv:1703.01245v1 [astro-ph.SR] 3 Mar 2017

6 March 2017

ABSTRACT

We present HAWK-I J-band light curves of five late-type T dwarfs (T6.5–T7.5) with a typical duration of four hours, and investigate the evidence for quasi-periodic photometric variability on intra-night timescales. Our photometry reaches precisions in the range 7–20 mmag, after removing instrumental systematics that correlate with sky background, seeing and airmass. Based upon a Lomb-Scargle periodogram analysis, the latest object in the sample - ULAS J2321 (T7.5) - appears to show quasi-periodic variability with a period of 1.64 hours and an amplitude of 3 mmag.

Given the low amplitude of variability and presence of systematics in our lightcurves, we discuss a Bayesian approach to robustly determine if quasi-periodic variability is present in a lightcurve affected by red noise. Using this approach, we conclude that the evidence for quasi-periodic variability in ULAS J2321 is not significant. As a result, we suggest that studies which identify quasi-periodic variables using the false alarm probability from a Lomb-Scargle periodogram are likely to overestimate the number of variable objects, even if field stars are used to set a higher false alarm probability threshold. Instead we argue that a hybrid approach combining a false alarm probability cut, followed by Bayesian model selection, is necessary for robust identification of quasi-periodic variability in lightcurves with red noise.

Key words: Stars: brown dwarfs

1 INTRODUCTION

Clouds form in the photospheres of brown dwarfs of most spectral types. L-type dwarfs have thick clouds of iron and silicates (e.g Tsuji et al. 1996; Allard et al. 2001; Marley et al. 2002; Burrows et al. 2006). Around the L/T transition these clouds either drop below the photosphere or break up into patches; as a result, the early T dwarfs are thought to be relatively cloud free (e.g. Ackerman & Marley 2001; Burgasser et al. 2002). In the late-T dwarfs condensates of alkali salts and sulfides are believed to form; models which include this opacity provide improved fits to the near- and mid-infrared colors of late-T dwarfs (Morley et al. 2012). However, it is worth noting that Line et al. (2015) find via hierarchical Bayesian model selection that (grey) clouds are not justified by the spectra of late-T dwarfs.

Observations of photometric variability in brown dwarfs provide a method for probing the presence and structure of condensate clouds. If the cloud deck is longitudinally heterogeneous, photometric variability will occur as parts of the cloud deck rotate in and out of view. Over the past 15 years, numerous surveys have attempted to detect photometric variability in brown dwarfs (see Metchev et al. 2015, and references within for a review). Major breakthroughs occurred when Artigau et al. (2009) detected strong ($\Delta J = 50$ mmag) quasi-periodic variability in the T2.5 dwarf SIMP J013656.5+093347 and Radigan et al. (2012) found even stronger ($\Delta J = 260$ mmag) quasi-periodic variability in the T1.5 dwarf 2MASS J21392676+0220226. These results triggered large surveys for near-infrared variability which showed that photometric variability is common, but typically of low amplitude. Radigan et al. (2014a) surveyed 57 L4–T9 dwarfs for variability; they found that 16% of their targets showed photometric variability in the J-band. Both the amplitude and frequency of variability is enhanced near the L/T transition, with 39% of L/T transition objects showing evidence for variability. Outside the transition typical amplitudes are 0.6–1.6%. The small amplitude of variability, and challenges of near-infrared photometry, mean that care must be taken to distinguish astrophysical variability from instrumental systematics. For example, the BAM survey (Wilson et al. 2014) found 14 variables amongst 69 surveyed objects, many with large amplitudes (>2%) and

© 0000 The Authors

at all spectral types. However a re-analysis of their data by Radigan (2014b) found that much of this variability was attributable to instrumental systematics. Space-based surveys can attain higher precision. Whilst they also suffer from instrumental systematics, in many cases these are well understood and can be removed from the data, to some extent. Metchev et al. (2015) surveyed 44 L3–T8 dwarfs with the Spitzer space telescope. They find that low-amplitude variability is exceedingly common, with 80% of L dwarfs varying with amplitudes $\geqslant 0.2\%$ and 40% of T dwarfs varying with amplitudes $\geqslant 0.4\%$.

These surveys have largely focused on spectral types of mid-T and earlier, typically including one or two later T dwarfs. Surveys for variability in late T dwarfs have the potential to test predictions of increased cloud opacity due to suflide and alkali salt condensates; the formation of new cloud layers may be accompanied by an increase in the frequency and amplitude of variability at later spectral types (Morley et al. 2014). The existing data show that variability exists amongst late T dwarfs; Metchev et al. (2015) finds clear variability in the T6 dwarf 2MASS J22282889-4310262 with an amplitude of 5% and a periodicity of 1.4 hours in the [3.6] band. Rajan et al. (2015) carried out small scale survey of three late-T dwarfs and one Y-dwarf. They report a detection of 13% variability with a periodicity of 3 hours in the T8.5 dwarf WISEP J045853.89+643452.9AB in one epoch, however this was not seen in a repeat visit. Recently, Cushing et al. (2016) find periodic variability with a period of 8.5 hours and semi-amplitude of 3.5% in the Y dwarf WISE J140518.39+553421.3. Whilst it is clear that variability exists amongst the late T and early Y dwarfs, more data is needed for good statistical estimates of its amplitude and frequency of occurrence.

As a caveat to the above discussion, we note that heterogeneous cloud coverage is not the only potential cause of variability in brown dwarfs. Magnetic spots are probably ruled out due to the neutrality of the atmosphere (Mohanty et al. 2002), although lightning discharge may increase the electron density in the atmosphere (Bailey et al. 2014). Another source of photometric variability may be temperature variations due to dynamical perturbations (Robinson & Marley 2014; Zhang & Showman 2014). Recently, photometric variability in a late M-dwarf has also been attributed to auroras (Hallinan et al. 2015). The photometric variability in this case is thought to be caused by non-thermal ionisation of the atmosphere produced by the impact of the auroral electron current. Similar photometric variability may be present in brown dwarfs of all spectral types; the coherent, pulsed radio emission associated with auroras is seen down to spectral types as late as T6.5 (Route & Wolszczan 2012, 2016). Indeed, it is worth pointing out that near-infrared variability and radio emission seem to be closely related, and that the prototype variable T-dwarf SIMP J013656.5+093347 has been detected in the radio (Kao et al. 2016). Hybrid variability mechanisms are also plausible, with either temperature fluctuations or non-thermal ionisation influencing the heterogeneity of the condensate cloud deck. Whilst it is likely that the majority of brown dwarf variability is due to condensate clouds, time resolved spectrophotometric observations are required to disentangle the contribution of other mechanisms.

Motivated by the relative sparsity of data on the variability of late T dwarfs, we undertook a mini-survey using HAWK-I on the VLT. We selected targets for observation from the ∼150 brown dwarfs with spectral types between T4 and T9 in the UKIDSS Large Area Survey (Burningham et al. 2013, 2010, and references therein). Many of these objects have multi-epoch observations as a result of their initial survey photometry and subsequent follow-up. We found that of the 95 objects with 2 epochs of J-band observations, 15% have discrepancies greater than 3σ in their photometry. To select the most likely candidates for photometric follow-up, we selected the 5 objects with spectral types close to T7 that have discrepant Y- and J-band photometry regardless of aperture size. These objects are listed in table 1.

The observations are described in section 2. The results are presented in section 3. In section 4 we describe our use of Bayesian model comparison to test for the presence of periodic variability. The results are discussed in section 5.

2 OBSERVATIONS AND DATA REDUCTION

On the nights between 17th Oct 2012 and 20th Oct 2012 we observed our target stars using HAWKI (Kissler-Patig et al. 2008) on UT4. Observing conditions are shown in table 1. All objects were placed near the centre of chip #3. Three objects (ULAS J2306, ULAS J2321 and WISE J2345) were observed in a 5-position offset pattern. 60 second exposures were taken in each offset position. We refer to this as pattern A. Two objects (ULAS J0150 and WISE J2340) were taken in a different offset pattern, with the aim of increasing time resolution. In this pattern 6 seperate 10 second exposures were taken at each offset position. We refer to this as pattern B.

The individual images were flat-fielded, dark- and background-subtracted using v.1.8.12 of the HAWKI pipeline recipes. The sky-background subraction is done in two steps. An initial pass is run on the first complete offset pattern (5 frames for pattern A, 30 frames for pattern B). In this pass a background image is produced from a median of the data and the data are background subtracted, aligned and co-added to produce a master frame in which objects are detected. The object mask is then used in a second pass to prevent bright objects from contaminating the background. For pattern A the background for an observation is the median of 7 frames either side of the observation. When computing the background value for each pixel, the two lowest and highest frames are rejected. For pattern B, 42 frames either side of the observation were used, and the five lowest and highest frames are rejected.

After background subtraction, the frames are aligned using the locations of bright stars and aperture photometry is carried out using the ULTRACAM data reduction software[1]. To achieve the best photometry we experimented with a range of aperture sizes and extraction methods. The best results were obtained with standard aperture photometry using apertures which were scaled with the seeing; aperture sizes ranged from 1 to 1.3 times the full-width-half-maximum (FWHM) measured in each frame.

We extracted photometry for our targets and all bright

[1] https://github.com/trmrsh/cpp-ultracam

Table 1. Journal of observations.

Name	Alias	SpT	UTC Start	UTC End	Seeing (″)	Airmass	Photometric?
WISE J234026.61-074508.1	WISE 2340-0745	T7	2012-10-17 23:40	2012-10-18 04:31	0.5–1.1	1.3 –1.7	N
ULAS J015024.37+135924.0	ULAS J0150+1359	T7.5	2012-10-18 04:39	2012-10-18 08:06	0.5–0.7	1.3–2.2	Y
WISE J234841.10-102844.1	WISE 2348-1028	T7	2012-10-18 23:35	2012-10-19 04:21	0.3–0.5	1.0–1.4	Y
ULAS J230601.02+130225.0	ULAS J2306+1302	T6.5	2012-10-19 23:34	2012-10-20 03:41	0.4–0.6	1.3–1.6	Y
ULAS J232123.79+135454.9	ULAS J2321+1354	T7.5	2012-10-20 23:31	2012-10-21 04:18	0.3–0.5	1.3–1.7	Y

stars located on the same chip as our target. The lightcurves of the bright stars were combined using a weighted mean to produce a reference lightcurve which was used to correct the lightcurve of our target star. We experimented with different choices of comparison stars to give the lightcurve with the smallest root-mean-square deviations from a constant flux. The results are shown in figure 1.

3 RESULTS

All of the brown dwarf lightcurves in figure 1 show smooth trends and higher frequency variability. The same patterns, with similar amplitudes, are seen in the lightcurves of our comparison stars. This strongly suggests that the variability is instrumental, rather than intrinsic in origin.

We can attempt to remove some of this instrumental variability by finding and removing correlations with airmass, sky background and seeing. Our data consists of a vector of times $\boldsymbol{t}$, fluxes $\boldsymbol{y}$ and uncertainties $\boldsymbol{\sigma}$. We also have measurements of the sky background $\boldsymbol{s}$, the average FWHM $\boldsymbol{f}$ and the airmass $\boldsymbol{X}$. For each object, the lower panel in figure 1 shows the data after dividing by $a\boldsymbol{s} + b\boldsymbol{f} + c\boldsymbol{X}$. The coefficients (a, b, c) are chosen to minimise the variance of the decorrelated data $\boldsymbol{y'}$. This process greatly reduces the variability apparent in the data, supporting our belief that much or all of this variability is due to instrumental systematics.

One possible source of instrumental variability is second-order extinction; in the J-band late T-dwarfs emit predominantly at wavelengths that are relatively unaffected by telluric absorption. Differences in telluric absorption between our targets and the references stars can introduce spurious variability. To some extent, this is accounted for by decorrelating agains airmass above. We also repeated the analysis above, adding a linear dependence on the relative humidity, which we use as a proxy for the line-of-sight water column. Adding this term does not significantly improve the removal of systematics.

For four of the five brown dwarfs, the residual variability is stochastic and has a typical amplitude of 10 mmag, or 1%. If intrinsic to the brown dwarfs, this may be caused by a turbulent atmosphere driving stochastic variability. However, we do not believe this variability is real, instead attributing it to instrumental systematics that have not been removed by the decorrelation process described above. The exception is the lightcurve of ULAS J2321+1354. The raw lightcurve appears to show a gradual linear trend and hints of periodic variability. Both the linear trend and the periodic variability survive the decorrelation process. We calculated a generalised Lomb-Scargle periodogram (Zechmeister & Kürster 2009), which accounts for a floating mean and observational errors. The periodogram was calculated using GATSPY (Vanderplas 2015), and shows a strong peak at a period of 1.64 hours. The periodogram, and best-fitting sinusoid are shown in figure 2. We calculated the false-alarm probability (FAP) for the strongest peak following Zechmeister & Kürster (2009). The FAP is given by

$$\mathrm{FAP} = 1 - [1 - \mathrm{Prob}(z > z_0)]^M, \qquad (1)$$

where $\mathrm{Prob}(z > z_0)$ is the probability that a periodogram power z can exceed the highest peak found z_0, under the assumption that the data are Gaussian noise. M is the number of independent frequencies. Using this formula we found FAP = 0.5%. Using the same method, the FAP for our comparison star lightcurves indicates that none of our comparison stars show significant evidence for periodicity; the lowest FAP amongst our comparison lightcurves was 10%.

Whilst the Lomb-Scargle periodogram of ULAS J2321+1354 shows evidence of periodic variability with a period of 1.64 hours, one cannot take the FAP at face value. Estimating the number of independent frequencies is difficult for unevenly spaced data. Monte-Carlo methods can overcome this (e.g Cumming et al. 1999), but are computationally expensive. In 100,000 trials where we calculated the periodogram of Gaussian noise with the same time sampling as the ULAS J2321+1354 lightcurve, 7 had a power greater than that seen in the actual periodogram, leading to an estimate of the false-alarm probability of $\mathrm{FAP_{MC}} = 0.007\%$.

Even Monte-Carlo methods cannot overcome a more serious issue in using the FAP to estimate the reality of an apparent periodic signal. Critically, the FAP gives the probability that a peak as large as the one observed will occur in the periodogram by chance, *under the assumption that the data are pure Gaussian noise*. Many astrophysical lightcurves do not satisfy this assumption. Instrumental systematics or astrophysical red noise mean that even in the absence of periodic signals, the lightcurve has clear temporal structure which is not consistent with pure Gaussian noise. This is clearly the case for our data, where instrumental systematics are present. As a result, we do not believe the FAP is an appropriate measure to judge the reality of the periodic signal in ULAS J2321+1354. Instead, we describe below an alternative method to determine the reality of a periodic signal in data which also shows systematics or astrophysical red noise, based on computing the Bayesian evidence for competing models.

4 BAYESIAN MODEL FITTING

Consider two models $\mathcal{M}_A$ and $\mathcal{M}_B$. Given our dataset $\boldsymbol{d}$ we can write the posterior odds ratio:

$$\frac{P(\mathcal{M}_A \mid \boldsymbol{d})}{P(\mathcal{M}_B \mid \boldsymbol{d})} = \frac{Pr(\mathcal{M}_A)}{Pr(\mathcal{M}_B)} \cdot \frac{m(\boldsymbol{d} \mid \mathcal{M}_A)}{m(\boldsymbol{d} \mid \mathcal{M}_A)}. \qquad (2)$$

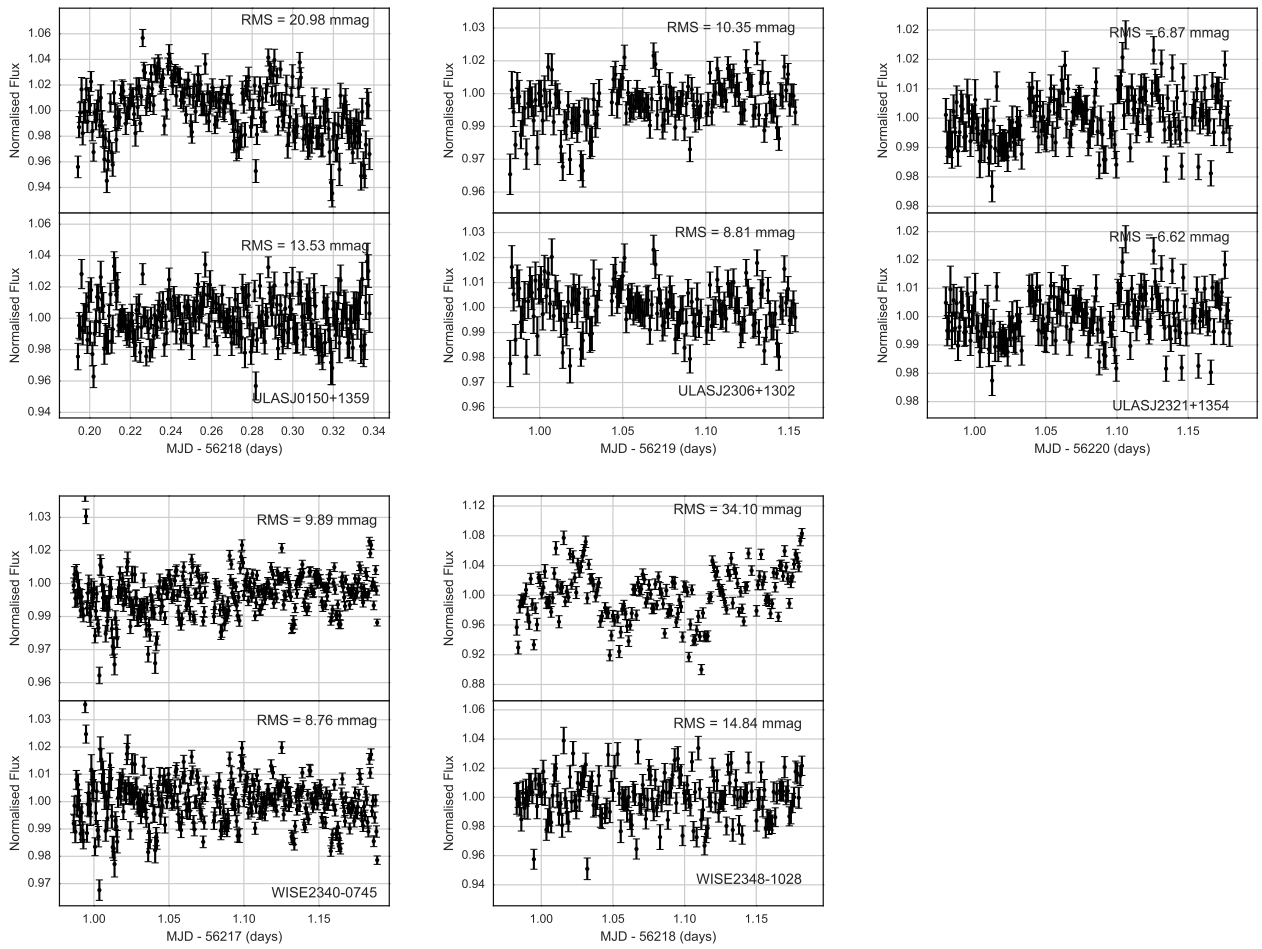

Figure 1. HAWKI *J*-band lightcurves for all stars observed. Each sub-panel shows the data before (top) and after (bottom) removing systematics (see text for details). The lightcurves have been normalised by dividing a constant fit to the data.

$\frac{Pr(\mathcal{M}_A)}{Pr(\mathcal{M}_B)}$ is the prior odds ratio: we set this ratio to unity as we have no prior reason to prefer either model. The second term on the right-hand side of the equation is the ratio of marginal likelihoods, also known as the evidence ratio or Bayes' factor. The marginal likelihood m of a dataset $\boldsymbol{d}$ given a model $\mathcal{M}_\mathcal{A}$ with parameter set θ_A is given by:

$$m(\boldsymbol{d} \mid \mathcal{M}_\mathcal{A}) = \int \mathcal{L}(\boldsymbol{d} \mid \mathcal{M}_\mathcal{A}, \theta_A)\, Pr(\theta_A \mid \mathcal{M}_\mathcal{A})\, d\theta_A, \quad (3)$$

where $\mathcal{L}(\boldsymbol{d} \mid \mathcal{M}_\mathcal{A}, \theta_A)$ is the likelihood function for the data, given the model and a particular set of parameters.

Under this framework, we could determine if a periodic signal is real by computing the odds ratio between two models. Our first model, $\mathcal{M}_\mathcal{A}$, would contain parameters to fit both the red noise in the dataset and a periodic component. The second model, $\mathcal{M}_\mathcal{B}$, would only contain the red noise parameters. A Bayesian approach to finding periodic signals thus has two challenges; finding appropriate models for the dataset, and adopting a suitable algorithm to calculate the marginal likelihood of these models, given these data.

4.1 Gaussian process models for red noise and periodicity

Gaussian processes (GPs) are seeing increasing use to model systematics and red noise in astrophysical datasets. For a textbook introduction see Rasmussen & Williams (2006). Roberts et al. (2013) offers a clear explanation of how GPs can be used to represent time-series data, and examples of applications to a wide range of datasets. Our dataset $\boldsymbol{d}$ consists of n points with times $\boldsymbol{t}$, fluxes $\boldsymbol{y}$ and uncertainties $\boldsymbol{\sigma}$. A GP represents the dataset as a multivariate Gaussian distribution. The covariance matrix $\boldsymbol{K}$ describes the extent to which each pair of data correlate with each other. Pure white noise would be described by covariance matrix where only the diagonal elements were non-zero. Allowing every element of the covariance matrix to be fitted would require $n \times n$ parameters, so the problem is made tractable by adopting a kernel function $k(t_i, t_j)$. The elements of the covariance matrix are then given by:

$$K_{ij} = \sigma_i^2 \delta_{ij} + k(t_i, t_j). \quad (4)$$

Various choices of kernel function exist; common choices for modelling red noise are the Matérn-3/2 (M32) function:

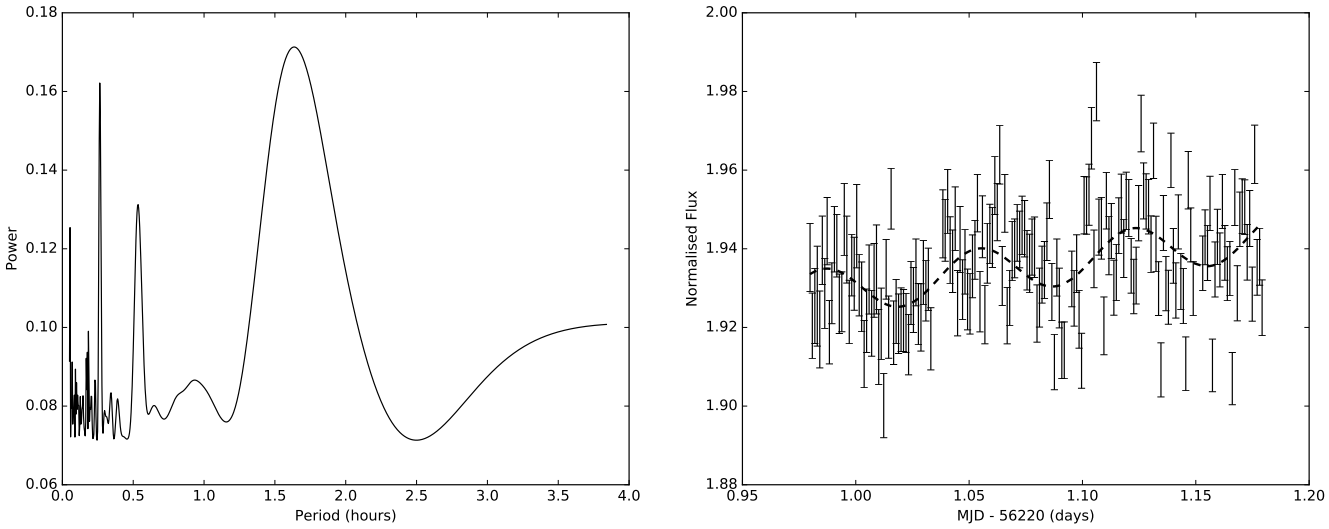

Figure 2. Left: The generalised Lomb-Scargle periodogram of the J-band lightcurve of ULAS J2321+1354. **Right:** The J-band lightcurve of ULAS J2321+1354 before decorrelation. The best fitting sinusoid and linear trend, with a period fixed at a value of 1.64 hours as suggested by the periodogram, is shown as a dashed line. The amplitude of the best fitting sinusoid is 3 mmag.

$$k(t_i, t_j) = a^2 \left(1 + \sqrt{\frac{3r^2}{\tau^2}}\right) \exp\left(-\sqrt{\frac{3r^2}{\tau^2}}\right), \quad (5)$$

and the squared-exponential (SE) function:

$$k(t_i, t_j) = a^2 \exp\left(-\frac{r^2}{2\tau^2}\right), \quad (6)$$

where $r^2 = (t_i - t_j)^2$. a and τ are the parameters of the GP model; a represents the typical amplitude of the variability due to red noise, and τ represents the timescale of variations. SE kernel functions give rather smooth variations, whereas M32 kernel functions are better suited to model rougher variations (Roberts et al. 2013).

The lightcurves shown in figure 1 appear to show systematics on two timescales; there are short timescale systematics which have quite rough structure and smoother, long timescale systematics. Motivated by this observation, we adopt the following kernel function to model the systematics in our data:

$$k(t_i, t_j) = a_s^2 \left(1 + \sqrt{\frac{3r^2}{\tau_s^2}}\right) \exp\left(-\sqrt{\frac{3r^2}{\tau_s^2}}\right) + a_l^2 \exp\left(-\frac{r^2}{2\tau_l^2}\right), \quad (7)$$

where a_s and τ_s are the amplitudes and timescales respectively of the short timescale variability and a_l and τ_l are the amplitudes and timescales of the longer timescale trends.

To calculate the marginal likelihood for a model we first need to be able to evaluate the likelihood for a specific set of parameters. Given a function $f_{\theta'}(\boldsymbol{t})$ that has parameters θ' and represents the mean of the data, we can calculate the residuals, $\boldsymbol{r} = \boldsymbol{y} - f_{\theta'}(\boldsymbol{t})$. The likelihood is given by (Rasmussen & Williams 2006):

$$\ln \mathcal{L}(\boldsymbol{d} \mid \theta) = -\frac{1}{2}\boldsymbol{r}^T \boldsymbol{K}^{-1} \boldsymbol{r} - \frac{1}{2} \ln \det \boldsymbol{K} - \frac{n}{2} \ln 2\pi, \quad (8)$$

where θ is the full set of parameters, including θ' and the parameters for the Gaussian process. We used the GEORGE[2] package (Ambikasaran et al. 2014) to implement our GP kernels.

Finally then, we are ready to define two models we can compare to test for the presence of periodic variability. Our first model, $\mathcal{M}_\mathcal{A}$, consists of the kernel function given by equation 7 and a sinusoidal mean function:

$$f_{\theta'}(\boldsymbol{t}) = \mu + A \sin\left[2\pi\left(\frac{\boldsymbol{t} - t_0}{P}\right)\right], \quad (9)$$

where μ is the mean level of the lightcurve, A the amplitude of the sinusoidal signal, t_0 represents the phase of the sinusoid and P is the rotational period of the brown dwarf. Our second model, $\mathcal{M}_\mathcal{B}$, combines the same kernel function with a simple mean function with one parameter:

$$f_{\theta'}(\boldsymbol{t}) = \mu. \quad (10)$$

These two models would be sufficient to test for the presence of a non-evolving, purely sinusoidal signal. However, the rotational variability seen in many brown dwarfs to date is more complex than this (e.g Metchev et al. 2015; Radigan et al. 2012; Artigau et al. 2009) . Multiple surface features, and evolution of the surface features on timescales comparable to the rotational period, can cause the lightcurve to differ drastically from a simple sinusoid. We describe a third model, $\mathcal{M}_\mathcal{C}$, which combines the simple mean function of equation 10 with a Gaussian process to represent

[2] http://dan.iel.fm/george

quasi-periodic variability and short timescale systematics. The kernel function for this Gaussian process is given by:

$$k_{\rm rot}(t_i,t_j) = A^2 \exp\left[-\Gamma \sin^2\left(\frac{\pi|r|}{P}\right) - \frac{r^2}{\tau_l^2}\right] + a_s^2\left(1+\sqrt{\frac{3r^2}{\tau_s^2}}\right)\exp\left(-\sqrt{\frac{3r^2}{\tau_s^2}}\right). \quad (11)$$

The first term in this kernel function is the quasi-periodic function used to model rotational variability by Aigrain et al. (2016), Vanderburg et al. (2015) and Haywood et al. (2014). This consists of an exponential-sine-squared kernel (ESn2) multiplied by an squared-exponential (SE) kernel. The addition of a M32 kernel models the short timescale systematics. P is the rotational period and A is the amplitude of quasi-periodic variability. The parameter Γ is a "roughness" parameter; it controls the regularity of the quasi-periodic lightcurve. Small values of Γ will give sinusoidal lightcurves, whereas large values will produce more irregular variations with larger contributions from harmonics. τ_l is the evolutionary timescale of the periodic variations.

The ability of this Gaussian process to represent quasi-periodic variability with additional systematics is demonstrated in figure 3. We created some fake data by adding together a sinusoidal signal, a longer term linear trend, white noise and red noise produced using a GP with a M32 kernel function. We then calculated the mean of a Gaussian process conditioned on the fake data for three different kernel functions. Figure 3 demonstrates that a Gaussian process using a ESn2 kernel function alone can represent periodic variability, but not a long term trend or any evolution in the variability. Multiplying the ESn2 kernel by a SE kernel gives the first term in equation 11, and can represent quasi-periodic variability and long term trends. The addition of the M32 kernel allows the Gaussian process to reproduce the quasi-periodic variability and the shorter timescale systematics. Model $\mathcal{M}_\mathcal{C}$ is thus capable of faithfully reproducing the lightcurves of variable brown dwarfs which are affected by instrumental systematics or astrophysical red noise (e.g. flaring).

4.2 Calculating the marginal likelihoods

To test if a periodic signal is present in a dataset, we need to be able to calculate the marginal likelihoods for our models $\mathcal{M}_\mathcal{A}$, $\mathcal{M}_\mathcal{B}$, $\mathcal{M}_\mathcal{C}$ along with the uncertainty on the marginal likelihoods. Direct integration of equation 3 over all of the allowed parameter space is computationally too expensive. Chib & Jeliazkov (2001) outline a method whereby the optimal parameters for a model $\mathcal{M}$ can be found via Markov-Chain Monte-Carlo (MCMC) techniques and the marginal likelihood can be estimated from the MCMC chains themselves. This requires a very large number of steps in the MCMC chains, since the marginal likelihood estimate will be dominated by a small fraction of steps which lie close to the maximum likelihood.

Instead, we use parallel-tempering MCMC (see Earl & Deem 2005, for a description of the algorithm) and calculate the marginal likelihood using thermodynamic integration (Goggans & Chi 2004). We use the parallel tempering algorithm as implemented in EMCEE (Foreman-Mackey et al. 2013). N parallel MCMC chains are run. Each chain samples from a modified posterior probability given by

$$\pi_T(\theta) = \mathcal{L}(\boldsymbol{d} \mid \theta)^{1/T} Pr(\theta), \quad (12)$$

where T is the temperature of the chain. For "hot" chains, the posterior becomes the prior and the chain explores a wide range of parameter space. Cold chains eplore the peaks of the likelihood function. During the MCMC fit, chains of different temperatures swap parameter values; this helps convergence if the likelihood function is multi-modal.

We use parallel-tempering MCMC not to improve convergence, but to allow an estimate of the marginal likelihood. We define the inverse temperature $\beta = 1/T$. The marginal likelihood as a function of inverse temperature is

$$m(\beta) = \int \mathcal{L}(\boldsymbol{d} \mid \theta)^\beta \, Pr(\theta)\, d\theta. \quad (13)$$

Goggans & Chi (2004) show that the marginal likelihood for a model can be written as:

$$\ln m = \int_0^1 \langle \ln \mathcal{L}(\boldsymbol{d} \mid \theta)^\beta \rangle d\beta, \quad (14)$$

where $\langle \ln \mathcal{L}(\boldsymbol{d} \mid \theta)^\beta \rangle$ is the average log-likelihood of an MCMC chain at inverse temperature β. This integral can easily be estimated using a quadrature formula from the parallel-tempering MCMC chains, and the uncertainty on the integral arising from using a finite number of temperatures estimated.

4.3 Tests on simulated data

To demonstrate the application of the method to data in which a clear sinusoidal signal is visible but which is also affected by systematics, we added a sinusoidal signal with period 1.66 hours and amplitude of 30 mmag to the lightcurve of ULAS J2321+1354. We then fit this fake data with models $\mathcal{M}_\mathcal{A}$, $\mathcal{M}_\mathcal{B}$, $\mathcal{M}_\mathcal{C}$ using the parallel-tempering MCMC method described above. Each model was fit using 30 chains of different temperatures. Each chain had an ensemble of 400 walkers to help with convergence and the MCMC was run for 1000 burn-in steps and 1000 production steps. Priors were chosen to be uninformative in most cases. The rotational period P, had a uniform prior which was fixed between 0.5 and 6 hours. To prevent degeneracy between the kernel function terms for short and long timescale variability, τ_l had a log-uniform prior between 2 hours and 1 day whilst τ_s had a log-uniform prior between 1 minute and 1 hour. We tested that the best fits and marginal likelihood had no significant dependence on the exact choice of priors. There was also no significant change in best fit or marginal likelihood for MCMC runs from different starting positions, or with increased numbers of temperatures, walkers or steps.

The fits to the data are shown in figure 4. Once the optimal parameters for the Gaussian process are known, samples can be drawn from the conditional probability for the value of the Gaussian process, given the observed data (Rasmussen & Williams 2006). The mean and standard deviation of these samples is calculated, and also shown in figure 4. The considerable flexibility of GPs to represent time-series data means that reasonable fits are obtained in each case. However, the posterior probability of the best fits are very different for each model. For model $\mathcal{M}_\mathcal{A}$, the sinusoidal

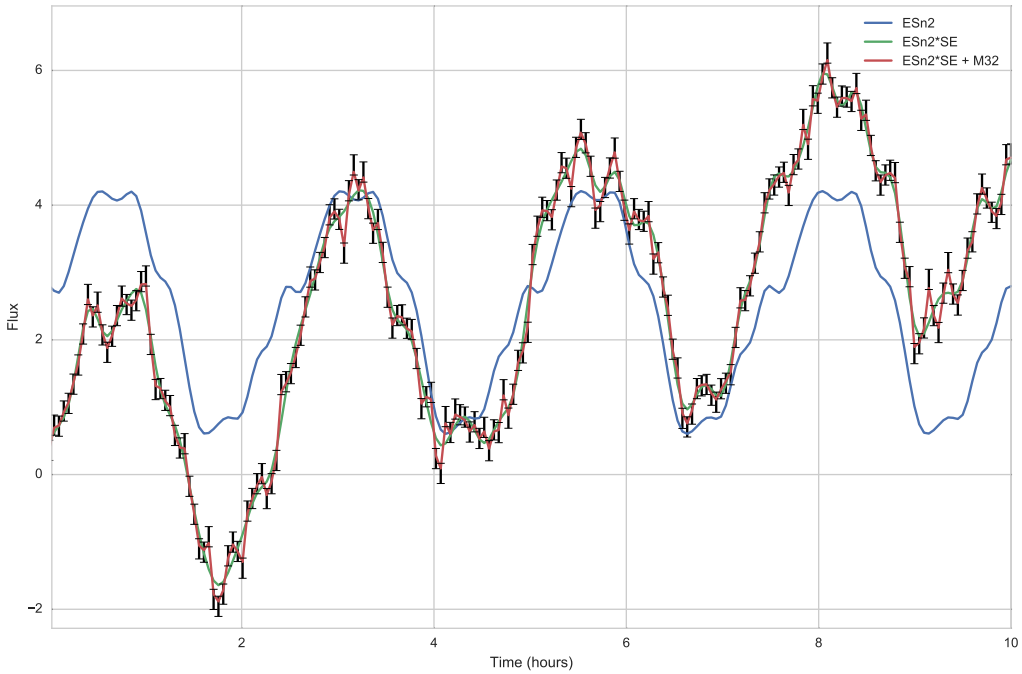

Figure 3. Representation of quasi-periodic variability with additional systematics by a Gaussian process. Fake data (black error bars) is created by adding a sinusoidal term, a long term linear trend, white noise, and red noise produced using a GP. Also plotted is the mean of Gaussian processes consisting of different kernel functions. An ESn2 kernel is shown in blue, the product of an ESn2 kernel and a SE kernel is shown in green, whilst the kernel function described in equation 11 is shown in red.

variability is modelled by the mean function, and the GP parameters are optimised to fit the systematics. For model $\mathcal{M}_\mathcal{B}$, the absence of a sinusoidal term means that a_l and τ_l are optimised to reproduce the periodic variability and the Gaussian process is not as effective at reproducing the systematics. This makes the posterior probability lower. This is reflected in the marginal likelihoods of the models. The marginal likelihoods for each model are $\ln m_\mathcal{A} = 628 \pm 3$, $\ln m_\mathcal{B} = 603 \pm 2$ and $\ln m_\mathcal{C} = 618 \pm 4$. Therefore both of the models with a periodic signal are clearly preferred to one with red noise alone, as expected.

4.4 ULAS J2321+1354

We applied the same model fitting process to the J-band lightcurve of ULAS J2321+1354. We do not attempt to remove systematics by detrending before fitting. To do so would mean that uncertainties in the detrending process are not reflected in the final parameter determinations. The fits to the lightcurve for each model, and the conditional distributions of the best fitting Gaussian process, are shown in figure 5. The Bayes' factors for each model show no preference for the models with a periodic component. The marginal likelihoods for each model are $\ln m_\mathcal{A} = 621 \pm 3$, $\ln m_\mathcal{B} = 620 \pm 2$ and $\ln m_\mathcal{C} = 621 \pm 3$.

One might argue that the instrumental systematics affecting the T-dwarfs and reference stars are similar. As discussed in section 3, secondary extinction may affect late T-dwarfs differently, but this is an exception: most systematic effects should influence target and reference stars alike. Thus, one may do better than adopting uninformative priors on the red noise terms in our models. Instead, we can fit model $\mathcal{A}$ to our reference lightcurves, and adopt the posterior distribution of parameters as our priors on the matching parameters in models $\mathcal{B}$ and $\mathcal{C}$ when fitting ULAS J2321+1354. However, this makes little difference to our results. Adopting this approach yields marginal likelihoods for each model of $\ln m_\mathcal{A} = 632 \pm 2$, $\ln m_\mathcal{B} = 633 \pm 2$ and $\ln m_\mathcal{C} = 633 \pm 3$.

4.5 Constraints on periodic variability

We attempted to place limits on the maximum amplitude of periodic variability in our target stars. To do so, we fit model $\mathcal{B}$ to the un-corrected data for each star and calculated the 99th-percentile limit on the amplitude of the periodic term as a function of period. The results are shown in figure 6. As one might expect, we are able to place tighter limits on the amplitude of periodic variability at shorter periods, and constraints at longer periods are weaker. Typically, we are able to set limits of 0.5–1.0 percent on periodic variability with timescales less than 5 hours. Constraints for ULASJ0150+1359 and WISEJ2348-1020 are weaker due to the larger amplitude of systematics in the lightcurves for these targets.

Since we have detected no evidence for short-term quasi-periodic variability in five targets that were selected on the basis of discrepant $Y-$ and $J-$band photometry on longer

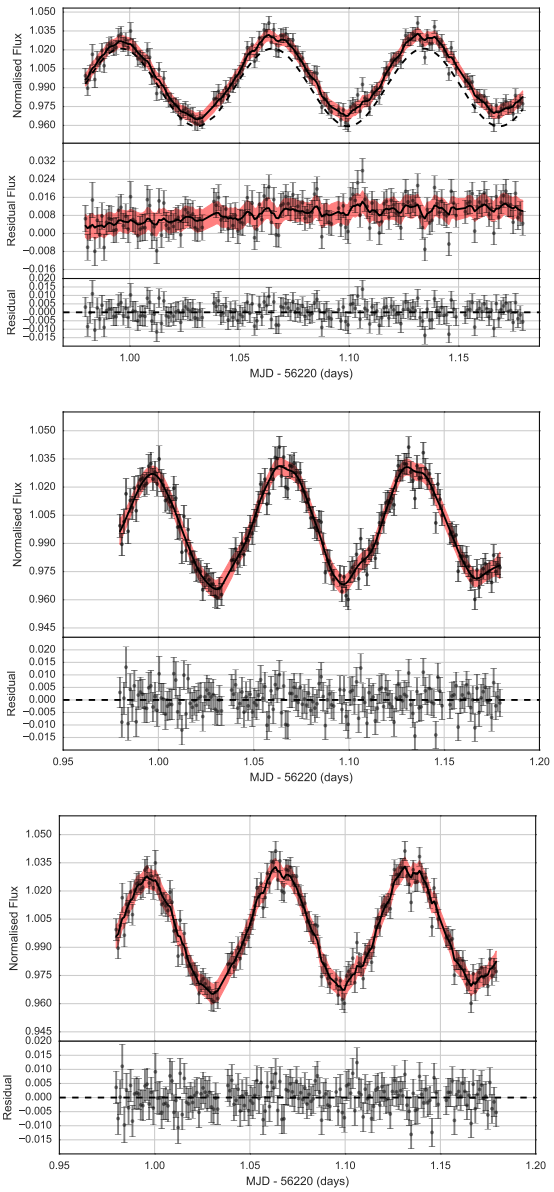

Figure 4. Fits to a simulated dataset (see text for details). From top to bottom we show fits using $\mathcal{M}_\mathcal{A}$ (red noise plus a sinusoid), $\mathcal{M}_\mathcal{B}$ (red noise alone) and $\mathcal{M}_\mathcal{C}$ (red noise plus quasi-periodic variability). In each plot the top panel shows the data and the mean (black line) and standard deviation (red region) of the best-fitting model. The bottom panel shows the residuals to the fit. For $\mathcal{M}_\mathcal{A}$ (top), a dashed line in the top panel shows the sinusoidal contribution to the best fit. The middle panel of this plot shows the data and the sinusoidal contribution subtracted, with the Gaussian process red noise model overplotted.

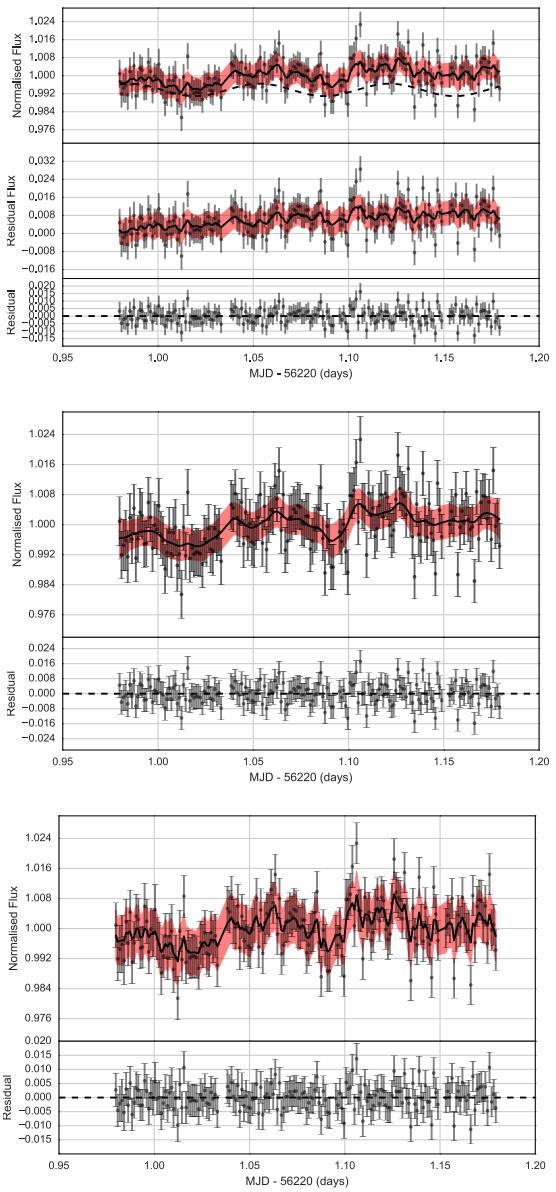

Figure 5. Fits the the J-band lightcurve of ULAS J2321+1354. From top to bottom we show fits using $\mathcal{M}_\mathcal{A}$ (red noise plus a sinusoid), $\mathcal{M}_\mathcal{B}$ (red noise alone) and $\mathcal{M}_\mathcal{C}$ (red noise plus quasi-periodic variability). In each plot the top panel shows the data and the mean (black line) and standard deviation (red region) of the best-fitting model. The bottom panel shows the residuals to the fit. For $\mathcal{M}_\mathcal{A}$ (top), a dashed line in the top panel shows the sinusoidal contribution to the best fit. The middle panel of this plot shows the data and the sinusoidal contribution subtracted with the Gaussian process red noise model overplotted.

timescales, we investigated if secondary extinction could be responsible for the discrepancies seen in the long-term photometry. By folding model atmospheres through the J-band filter, after applying telluric absorption using the Mauna Kea transmission profiles of Lord et al. (1992), we can estimate the impact of secondary extinction on our measurements for a range of precipitable water columns. We used COND model atmospheres (Allard et al. 2001) with effective temperatures of 1000 K and 4000 K to represent our target and a 'typical' reference star respectively. We found that changes in water column from 1 to 5 mm can introduce spurious variability of up to 0.1 mags if not properly accounted for. As a result, we cannot rule out the possibility that this is the origin of the discrepant $Y-$ and $J-$band photometry used to select our target stars.

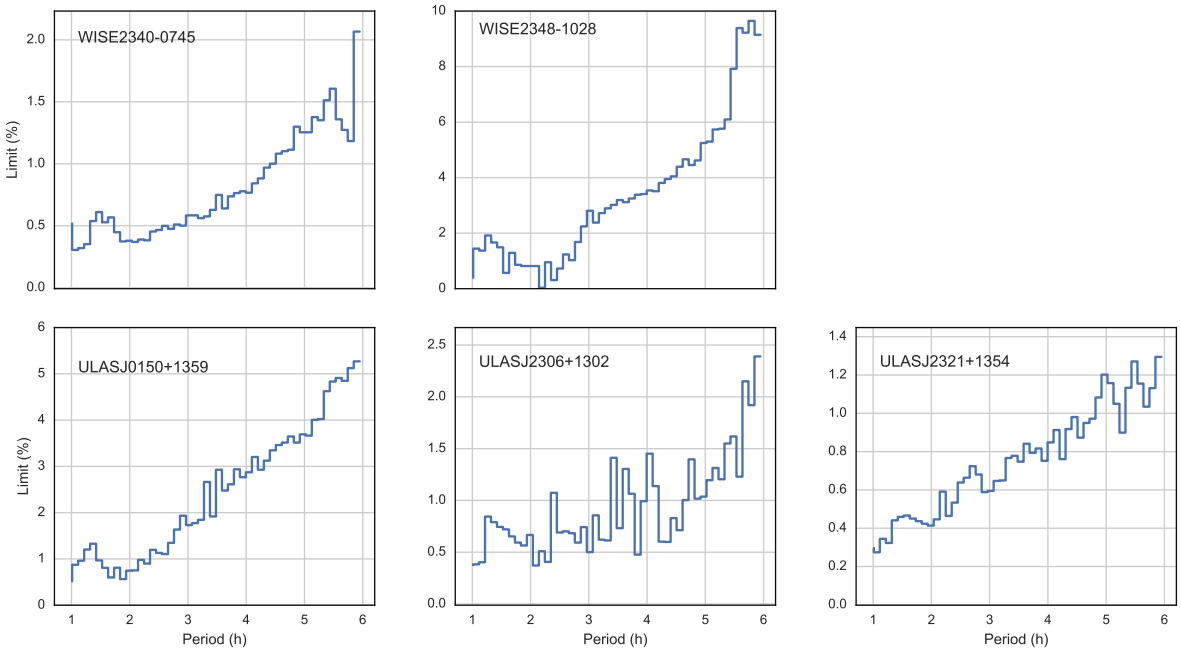

Figure 6. Constraints on periodic variability in our target stars. The lines show the 99% confidence limit to the amplitude of sinusoidal variability, as a function of period, for each star.

5 DISCUSSION

By comparing the evidence for models with and without periodic signals we see that there is no robust evidence for periodic variability in ULAS J2321+1354. This is in stark contrast to an analysis using Lomb-Scargle periodograms and the FAP. This is perhaps unsurprising, since the FAP measures the probability a given lightcurve could arise if the data were Gaussian (white) noise and the lightcurve of ULAS J2321+1354 shows clear red noise which we attribute to instrumental systematics.

The difficulty of interpreting FAPs in data where any sinusoidal signal is accompanied by red noise (either from instrumental systematics or astrophysical sources) is well known. A commonly adopted solution (see Littlefair et al. 2010; Metchev et al. 2015, for example) is to apply a stricter FAP criterion, where the critical FAP is arrived at via some ad-hoc approach. Usually this involves looking at the FAP distribution of stars assumed not to show periodic variability, and choosing a critical FAP high enough to avoid false positive claims. However, this approach is not ideal, since it assumes that the red noise in the control stars used is similar to that in the suspected periodic variables. This may well not be true. Had we adopted this approach for ULAS J2321+1354 we would have concluded that the periodic variability is real; it has by far the lowest FAP of all of the stars in the field of view.

Our approach using Bayesian model comparison reaches a different conclusion. This is because we are asking if the variability exhibited by ULAS J2321+1354 is likely to have arisen from systematics alone. Given the presence of both short timescale systematics and a long term systematic trend of similar amplitude to any possible periodic variability, it is not surprising that we find no evidence for a periodic signal in ULAS J2321+1354. The Bayesian model comparison approach is thus more conservative, and less likely to mistake instrumental systematics for quasi-periodic variability. It is, however, more computationally expensive, with our models taking ∼1 hour to run on a 64-core desktop machine. We would thus recommend a hybrid approach when dealing with large samples of stars, where candidates for periodic variability are first selected via a Lomb-Scargle periodogram and the reality of that signal is then assessed using model comparison.

6 CONCLUSIONS

We observed five late T dwarfs with HAWK-I on the VLT for a typical duration of ∼4 hours. Our data are affected by instrumental systematics with amplitudes between 7 and 30 mmags. Linear detrending for systematics that correlate with sky background, FWHM and airmass removes some, but not all, of these systematics. We are able to place limits on the amplitude of any short periodic variability of 0.5–5 percent, depending on the level of systematics in each lightcurve.

One object, ULAS J2321+13 appears to show periodic variability with a period of 1.64 hours and an amplitude of 3 mmags. A Lomb-Scargle periodogram appears to confirm this variability but a Bayesian comparison of models with and without periodic terms finds no significant evidence that this periodic signal is real. We therefore conclude that none of our targets show periodic variability with amplitudes of a 1.5 percent or greater.

7 ACKNOWLEDGEMENTS

SPL is supported by STFC grant ST/M001350/1. BB acknowledges financial support from the European Commision in the form of a Marie Curie International Outgoing Fellowship (PIOF-GA-2013-629435). Based on observations made with ESO Telescopes at the La Silla Paranal Observatory under programme ID 090.C-0721. ChH highlights financial support of the European Community under the FP7 by the ERC starting grant 257431. This research made use of Astropy, a community-developed core Python package for Astronomy (Astropy Collaboration, 2013). We thank the anonymous referee for their comments, which significantly improved the paper.

REFERENCES

Ackerman A. S., Marley M. S., 2001, ApJ, 556, 872
Aigrain S., Parviainen H., Pope B. J. S., 2016, MNRAS,
Allard F., Hauschildt P. H., Alexander D. R., Tamanai A., Schweitzer A., 2001, ApJ, 556, 357
Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O'Neil M., 2014, preprint, (arXiv:1403.6015)
Artigau É., Bouchard S., Doyon R., Lafrenière D., 2009, ApJ, 701, 1534
Bailey R. L., Helling C., Hodosán G., Bilger C., Stark C. R., 2014, ApJ, 784, 43
Burgasser A. J., Marley M. S., Ackerman A. S., Saumon D., Lodders K., Dahn C. C., Harris H. C., Kirkpatrick J. D., 2002, ApJ, 571, L151
Burningham B., et al., 2010, MNRAS, 406, 1885
Burningham B., et al., 2013, MNRAS, 433, 457
Burrows A., Sudarsky D., Hubeny I., 2006, ApJ, 640, 1063
Chib S., Jeliazkov I., 2001, Journal of the American Statistical Association, 96, 270
Cumming A., Marcy G. W., Butler R. P., 1999, ApJ, 526, 890
Cushing M. C., et al., 2016, preprint, (arXiv:1602.06321)
Earl D. J., Deem M. W., 2005, Physical Chemistry Chemical Physics (Incorporating Faraday Transactions), 7, 3910
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306
Goggans P. M., Chi Y., 2004, in Erickson G. J., Zhai Y., eds, American Institute of Physics Conference Series Vol. 707, Bayesian Inference and Maximum Entropy Methods in Science and Engineering. pp 59–66, doi:10.1063/1.1751356
Hallinan G., et al., 2015, Nature, 523, 568
Haywood R. D., et al., 2014, MNRAS, 443, 2517
Kao M. M., Hallinan G., Pineda J. S., Escala I., Burgasser A., Bourke S., Stevenson D., 2016, ApJ, 818, 24
Kissler-Patig M., et al., 2008, A&A, 491, 941
Line M. R., Teske J., Burningham B., Fortney J. J., Marley M. S., 2015, ApJ, 807, 183
Littlefair S. P., Naylor T., Mayne N. J., Saunders E. S., Jeffries R. D., 2010, MNRAS, 403, 545
Lord S. D., Hollenbach D. J., Colgan S. W. J., Haas M. R., Rubin R. H., Erickson E. F., Carral P., Maloney P., 1992, in American Astronomical Society Meeting Abstracts. p. 1182
Marley M. S., Seager S., Saumon D., Lodders K., Ackerman A. S., Freedman R. S., Fan X., 2002, ApJ, 568, 335
Metchev S. A., et al., 2015, ApJ, 799, 154
Mohanty S., Basri G., Shu F., Allard F., Chabrier G., 2002, ApJ, 571, 469
Morley C. V., Fortney J. J., Marley M. S., Visscher C., Saumon D., Leggett S. K., 2012, ApJ, 756, 172
Morley C. V., Marley M. S., Fortney J. J., Lupu R., 2014, ApJ, 789, L14
Radigan J., 2014b, ApJ, 797, 120
Radigan J., Jayawardhana R., Lafrenière D., Artigau É., Marley M., Saumon D., 2012, ApJ, 750, 105
Radigan J., Lafrenière D., Jayawardhana R., Artigau E., 2014a, ApJ, 793, 75
Rajan A., et al., 2015, MNRAS, 448, 3775
Rasmussen C. E., Williams C. K. I., 2006, Gaussian processes for machine learning. Adaptive Computation and Machine Learning, MIT Press, Cambridge, MA, http://www.ams.org/mathscinet-getitem?mr=MR2514435
Roberts S., Osborne M., Ebden M., Reece S., Gibson N., Aigrain S., 2013, Philosophical Transactions: Mathematical, Physical and Engineering Sciences, 371, 1
Robinson T. D., Marley M. S., 2014, ApJ, 785, 158
Route M., Wolszczan A., 2012, ApJ, 747, L22
Route M., Wolszczan A., 2016, ApJ, 821, L21
Tsuji T., Ohnaka K., Aoki W., Nakajima T., 1996, A&A, 308, L29
Vanderburg A., et al., 2015, ApJ, 800, 59
Vanderplas J., 2015, gatspy: General tools for Astronomical Time Series in Python, doi:10.5281/zenodo.14833, http://dx.doi.org/10.5281/zenodo.14833
Wilson P. A., Rajan A., Patience J., 2014, A&A, 566, A111
Zechmeister M., Kürster M., 2009, A&A, 496, 577
Zhang X., Showman A. P., 2014, ApJ, 788, L6